arXiv:1609.09726v1 [cs.NI] 30 Sep 2016

Revisiting 802.11 Rate Adaptation from Energy Consumption's Perspective

Iñaki Ucar*, Carlos Donato†, Pablo Serrano*, Andres Garcia-Saavedra‡, Arturo Azcorra*†, Albert Banchs*†
inaki.ucar@uc3m.es, {pablo,azcorra,banchs}@it.uc3m.es, carlos.donato@imdea.org, andres.garcia.saavedra@neclab.eu

*Universidad Carlos III de Madrid, Avda. Universidad, 30, 28911 Leganés, Spain
†IMDEA Networks Institute, Avda. Mar Mediterráneo, 22, 28918 Leganés, Spain
‡NEC Labs Europe, Kurfürsten-Anlage, 36, 69115 Heidelberg, Germany

ABSTRACT

Rate adaptation in 802.11 WLANs has received a lot of attention from the research community, with most of the proposals aiming at maximising throughput based on network conditions. Considering energy consumption, an implicit assumption is that optimality in throughput implies optimality in energy efficiency, but this assumption has been recently put into question. In this paper, we address via analysis and experimentation the relation between throughput performance and energy efficiency in multi-rate 802.11 scenarios. We demonstrate the trade-off between these performance figures, confirming that they may not be simultaneously optimised, and analyse their sensitivity towards the energy consumption parameters of the device. Our results provide the means to design novel rate adaptation schemes that takes energy consumption into account.

Keywords

WLAN; 802.11; Rate Adaptation; Energy Efficiency

1. INTRODUCTION

In recent years, along with the exponential growth in mobile data applications and the corresponding traffic volume demand (see e.g. [1]), we have witnessed an increased attention towards "green operation" of networks, which is required to support a sustainable growth of the communication infrastructures. For the case of wireless communications, there is the added motivation of a limited energy supply (i.e., batteries), which has triggered a relatively large amount of work on energy efficiency [9]. It turns out, though, that energy efficiency and performance do not necessarily come hand in hand, as some recent research has pointed out [2,3], and that a criterion may be required to set a proper balance between them.

This paper is devoted to the problem of rate adaptation (RA) in 802.11 WLANs from the energy consumption's perspective. RA algorithms are responsible for selecting the most appropriate modulation and coding scheme (MCS) and transmission power (TXP) to use, given an estimation of the link conditions, and have received a vast amount of attention from the research community (see e.g. [4] and references therein). In general, the challenge lies in distinguishing between those loses due to collisions and those due to poor radio conditions, because they should trigger different reactions. In addition, the performance figure to optimise is commonly the throughput or a related one such as, e.g., the time required to deliver a frame.

It is generally assumed that optimality in terms of throughput also implies optimality in terms of energy efficiency. However, some recent work [6,7] has shown that throughput maximisation does not result in energy efficiency maximisation, at least for 802.11n. However, we still lack a proper understanding of the causes behind this "non-duality", as it may be caused by the specific design of the algorithms studied, the extra consumption caused by the complexity of MIMO techniques, or any other reason. In fact, it could be an inherent trade-off given by the power consumption characteristics of 802.11 interfaces, and, if so, RA techniques should not be agnostic to this case.

This work tackles the latter question from a formal standpoint. A question which, to the best of the authors' knowledge, has never been addressed in the literature. For this purpose, and with the aim of isolating the variables of interest, we present a joint goodput and energy consumption model for single 802.11 spatial streams in the absence of interfering traffic. Packet losses occur due to poor channel conditions and RA can tune only two variables: MCS and TXP.

Building on this model, we provide the following contributions: (*i*) we demonstrate through an extensive numerical evaluation that energy consumption and throughput performance are different optimisation objectives in 802.11, and not only an effect of MIMO or certain algorithms' suboptimalities; (*ii*) we analyse the relative impact of each energy consumption component on the resulting performance of RA, which serves to identify the critical factors to consider for the design of RA algorithms, and illustrate that different hardware should employ different configurations; and (*iii*) we experimentally validate our numerical results.

The rest of this paper is organised as follows. In Section 2, we develop the theoretical framework: a joint goodput-energy model built around separate previous models. In Section 3, we provide a detailed analysis of the trade-off between energy efficiency and maximum goodput, including a discus-

©2016 Copyright held by the authors. This is the authors' version of the work. It is posted here for your personal use. Not for redistribution. The definitive version was published in

ACM MSWiM '16, November 13 - 17, 2016, Malta, Malta

ACM ISBN 978-1-4503-4502-6/16/11.
DOI: http://dx.doi.org/10.1145/2988287.2989149

sion of the role of the different energy parameters involved. In Section 4, we support our numerical analysis with experimental results. Finally, Section 5 summarises the paper.

2. JOINT GOODPUT-ENERGY MODEL

In this section, we develop a joint goodput-energy model for a single 802.11 spatial stream and the absence of interfering traffic. It is based on previous studies about goodput and energy consumption of wireless devices. As stated in the introduction, the aim of this model is the isolation of the relevant variables (MCS and TXP) to let us delve in the relationship between goodput and energy consumption optimality in the absence of other effects such as collisions or MIMO.

Beyond this primary intent, it is worth noting that these assumptions conform with real-world scenarios in the scope of recent trends in the IEEE 802.11 standard development, namely, the amendments 11ac and 11ad, where device-to-device communications (mainly through beamforming and MU-MIMO) are of paramount importance.

2.1 Goodput Model

We base our study on the work by Qiao *et al.* [8], which develops a robust goodput model that meets the established requirements. This model analyses the IEEE 802.11a Distributed Coordination Function (DCF) over the assumption of an AWGN (Additive White Gaussian Noise) channel without interfering traffic.

Let us briefly introduce the reader to the main concepts, essential to our analysis, of the goodput model by Qiao *et al.*. Given a packet of length l ready to be sent, a frame retry limit $n_{\max}$ and a set of channel conditions $\hat{s} = \{s_1, \ldots, s_{n_{\max}}\}$ and modulations $\hat{m} = \{m_1, \ldots, m_{n_{\max}}\}$ used during the potential transmission attempts, the *expected effective goodput* $\mathcal{G}$ is modelled as the ratio between the expected delivered data payload and the expected transmission time as follows:

$$\mathcal{G}(l, \hat{s}, \hat{m}) = \frac{\mathbb{E}\left[\text{data}\right]}{\mathbb{E}\left[\mathcal{D}_{\text{data}}\right]} = \frac{\Pr[\text{succ} \mid l, \hat{s}, \hat{m}] \cdot l}{\mathbb{E}\left[\mathcal{D}_{\text{data}}\right]} \tag{1}$$

where $\Pr[\text{succ} \mid l, \hat{s}, \hat{m}]$ is the probability of successful transmission conditioned to $l, \hat{s}, \hat{m}$, given by Equation (5) in [8]. The expected transmission time is defined as follows:

$$\begin{aligned}\mathbb{E}\left[\mathcal{D}_{\text{data}}\right] = {} & (1 - \Pr[\text{succ} \mid l, \hat{s}, \hat{m}]) \cdot \mathcal{D}_{\text{fail}|l,\hat{s},\hat{m}} \\ & + \Pr[\text{succ} \mid l, \hat{s}, \hat{m}] \cdot \mathcal{D}_{\text{succ}|l,\hat{s},\hat{m}}\end{aligned} \tag{2}$$

where

$$\begin{aligned}\mathcal{D}_{\text{succ}|l,\hat{s},\hat{m}} = \sum_{n=1}^{n_{\max}} \Pr[n \text{ succ} \mid l, \hat{s}, \hat{m}] \cdot \Bigg\{ & \sum_{i=2}^{n_{\max}} \left[\overline{T}_{\text{bkoff}}(i)\right. \\ & \left. + T_{\text{data}}(l, m_i) + \overline{\mathcal{D}}_{\text{wait}}(i)\right] \\ & + \overline{T}_{\text{bkoff}}(1) + T_{\text{data}}(l, m_1) + T_{\text{SIFS}} \\ & + T_{\text{ACK}}(m'_n) + T_{\text{DIFS}} \Bigg\}\end{aligned} \tag{3}$$

is the average duration of a successful transmission and

$$\begin{aligned}\mathcal{D}_{\text{fail}|l,\hat{s},\hat{m}} = \sum_{i=1}^{n_{\max}} & \left[\overline{T}_{\text{bkoff}}(i)\right. \\ & \left. + T_{\text{data}}(l, m_i) + \overline{\mathcal{D}}_{\text{wait}}(i+1)\right]\end{aligned} \tag{4}$$

is the average time wasted during the $n_{\max}$ attempts when the transmission fails.

$\Pr[n \text{ succ} \mid l, \hat{s}, \hat{m}]$ is the probability of successful transmission at the n-th attempt conditioned to $l, \hat{s}, \hat{m}$, and $\overline{\mathcal{D}}_{\text{wait}}(i)$ is the average waiting time before the i-th attempt. Their expressions are given by Equations (7) and (8) in [8]. The transmission time (T_{data}), ACK time (T_{ACK}) and average backoff time ($\overline{T}_{\text{bkoff}}$) are given by Equations (1), (2) and (3) in [8]. Finally, T_{SIFS} and T_{DIFS} are 802.11a parameters, and they can be found also in Table 2 in [8].

2.2 Energy Consumption Model

The selected energy model is our previous work of [10], which has been further validated via ad-hoc circuitry and specialised hardware [11] and, to the best of our knowledge, stands as the most accurate energy model for 802.11 devices published so far, because it accounts not only the energy consumed by the wireless card, but the consumption of the whole device. While classical models focused on the wireless interface solely, this one demonstrates empirically that the energy consumed by the device itself cannot be neglected as a device-dependent constant contribution. Conversely, devices incur an energy cost derived from the frame processing, which may impact the relationship that we want to evaluate in this paper.

This model can be summarised as follows:

$$\overline{P} = \rho_{\text{id}} + \sum_{i \in \{\text{tx},\text{rx}\}} \rho_i \tau_i + \sum_{i \in \{\text{g},\text{r}\}} \gamma_{\text{x}i} \lambda_i \tag{5}$$

where $\rho_{\text{id}}, \rho_{\text{tx}}, \rho_{\text{rx}}$ are the power consumed by the device in idle, transmission and reception states respectively; $\tau_{\text{tx}}, \tau_{\text{rx}}$ are the airtime percentages in transmission and reception; $\gamma_{\text{xg}}, \gamma_{\text{xr}}$ are the so called *cross-factors*, a per-frame energy toll for generation and reception respectively; and $\lambda_{\text{g}}, \lambda_{\text{r}}$ are the frame generation and reception rates.

Therefore, the average power consumed $\overline{P}$ is a function of five device-dependent parameters ($\rho_i, \gamma_{\text{x}i}$) and four traffic-dependent ones (τ_i, λ_i).

2.3 Energy Efficiency Analysis

Putting together both models, we are now in a position to build a joint goodput-energy model for 802.11a DCF. Let us consider the average durations (3) and (4). Based on their expressions, we multiply the idle time ($\overline{\mathcal{D}}_{\text{wait}}$, $\overline{T}_{\text{bkoff}}$, T_{SIFS}, T_{DIFS}) by ρ_{id}, the transmission time (T_{data}) by ρ_{tx}, and the reception time (T_{ACK}) by ρ_{rx}. The resulting expressions are the average energy consumed in a successful transmission $\mathcal{E}_{\text{succ}|l,\hat{s},\hat{m}}$ and the average energy wasted when a transmission fails $\mathcal{E}_{\text{fail}|l,\hat{s},\hat{m}}$:

$$\begin{aligned}\mathcal{E}_{\text{succ}|l,\hat{s},\hat{m}} = \sum_{n=1}^{n_{\max}} \Pr[n \text{ succ} \mid l, \hat{s}, \hat{m}] \cdot \Bigg\{ & \sum_{i=2}^{n_{\max}} \left[\rho_{\text{id}} \overline{T}_{\text{bkoff}}(i)\right. \\ & \left. + \rho_{\text{tx}} T_{\text{data}}(l, m_i) + \rho_{\text{id}} \overline{\mathcal{D}}_{\text{wait}}(i)\right] \\ & + \rho_{\text{id}} \overline{T}_{\text{bkoff}}(1) + \rho_{\text{tx}} T_{\text{data}}(l, m_1) + \rho_{\text{id}} T_{\text{SIFS}} \\ & + \rho_{\text{rx}} T_{\text{ACK}}(m'_n) + \rho_{\text{id}} T_{\text{DIFS}} \Bigg\}\end{aligned} \tag{6}$$

$$\begin{aligned}\mathcal{E}_{\text{fail}|l,\hat{s},\hat{m}} = \sum_{i=1}^{n_{\max}} & \left[\rho_{\text{id}} \overline{T}_{\text{bkoff}}(i)\right. \\ & \left. + \rho_{\text{tx}} T_{\text{data}}(l, m_i) + \rho_{\text{id}} \overline{\mathcal{D}}_{\text{wait}}(i+1)\right]\end{aligned} \tag{7}$$

Then, by analogy with (2), the *expected energy consumed*

Table 1: Modes of the IEEE 802.11a PHY

Mode Index	1	2	3	4	5	6	7	8
MCS (Mbps)	6	9	12	18	24	36	48	54

per frame transmitted, $\mathbb{E}[\mathcal{E}_{\text{data}}]$, can be written as follows:

$$\mathbb{E}[\mathcal{E}_{\text{data}}] = \gamma_{\text{xg}} + (1 - \Pr[\text{succ} \mid l, \hat{s}, \hat{m}]) \cdot \mathcal{E}_{\text{fail}|l,\hat{s},\hat{m}} + \Pr[\text{succ} \mid l, \hat{s}, \hat{m}] \cdot \mathcal{E}_{\text{succ}|l,\hat{s},\hat{m}} \quad (8)$$

It is noteworthy that the receiving cross-factor does not appear in this expression because ACKs (acknowledgements) are processed in the network card exclusively, and thus its processing toll is negligible.

Finally, we define the *expected effective energy efficiency* μ as the ratio between the expected delivered data payload and the expected energy consumed per frame, which can be expressed in *bits per Joule* (bpJ):

$$\mu(l, \hat{s}, \hat{m}) = \frac{\mathbb{E}[\text{data}]}{\mathbb{E}[\mathcal{E}_{\text{data}}]} \quad (9)$$

3. NUMERICAL RESULTS

Building on the joint model presented in the previous section, here we explore the relationship between optimal goodput and energy efficiency in 802.11a. More specifically, our objective is to understand the behaviour of the energy efficiency of a single spatial stream as the MCS and TXP change following our model to meet the optimal goodput.

3.1 Optimal Goodput

We note that the main goal of RA, generally, is to maximise the effective goodput that a station can achieve by varying the parameters of the interface. In terms of the model discussed in the previous section, a rate adaptation algorithm would aspire to fit the following curve:

$$\max \mathcal{G}(l, \hat{s}, \hat{m}) \quad (10)$$

We provide the numerical results for this goodput maximisation problem in Fig. 1, which are in good agreement with those obtained in [8]. For the sake of simplicity but without loss of generality we fix $l = 1500$ octets and $n_{\max} = 7$ retries, and assume that the channel conditions and the transmission strategy are constant across retries ($\hat{s} = \{s_1, \ldots, s_1\}$ and $\hat{m} = \{m_1, \ldots, m_1\}$).

Fig. 1 illustrates which mode (see Table 1) is optimal in terms of goodput, given an SNR level. We next address the question of whether this optimisation is aligned with energy efficiency maximisation.

3.2 Extension of the Energy Parametrisation

The next step is to delve into the energy consumption of wireless devices. [10] provides real measurements for five devices: three AP-like platforms (Linksys WRT54G, Raspberry Pi and Soekris net4826-48) and two hand-held devices (HTC Legend and Samsung Galaxy Note 10.1). Two of the four parameters needed are constant ($\rho_{\text{id}}, \gamma_{\text{xg}}$), and the other two ($\rho_{\text{tx}}, \rho_{\text{rx}}$) depend on the MCS and the TXP used. However, the characterisation done in [10] is performed for a subset of the MCS and TXP available, so we next detail how we extend the model to account for a larger set of operation parameters.

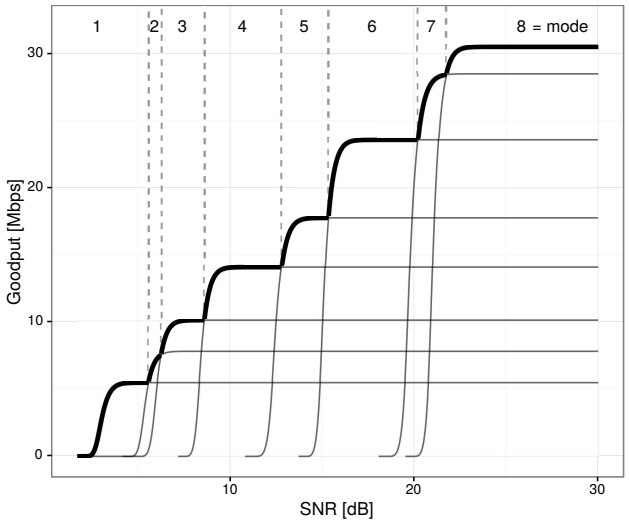

Figure 1: Optimal goodput (bold envelope) as a function of SNR.

A detailed analysis of the numerical figures presented in [10] suggests that ρ_{rx} depends linearly on the MCS, and that ρ_{tx} depends linearly on the MCS and the TXP (in mW). Based on these observations, we define the following linear models:

$$\rho_{\text{tx}} = \alpha_0 + \alpha_1 \cdot \text{MCS} + \alpha_2 \cdot \text{TXP} \quad (11)$$

$$\rho_{\text{rx}} = \beta_0 + \beta_1 \cdot \text{MCS} \quad (12)$$

The models are fed with the data reported in [10], and the resulting fitting is illustrated in Figs. 2a and 2b, while Table 2 collects the model estimates for each device (with errors between parentheses), as well as the adjusted r-squared. Since these linear models show a very good fit, they support the generation of synthetic data for the different MCS and TXP required.

3.3 Energy Consumption

To compute the energy consumption using the above parametrisation, first we have to define the assumptions for the considered scenario. We assume for simplicity a device-to-device communication, with fixed and reciprocal channel conditions during a sufficient period of time (i.e., low or no mobility). As we have discussed before, our primary goal is to isolate MCS and TXP as variables of interest, but we must not forget that these are also reasonable assumptions in scenarios targeted by recent 802.11 standard developments (11ac, 11ad).

For instance, given channel state information from a receiver, the transmitter may decide to increase the TXP in order to increase the receiver's SNR (and thus the expected goodput), or to decrease it if the channel quality is high enough. Although the actual relationship between TXP and SNR depends on the specific channel model (e.g., distance, obstacles, noise), without loss of generality, we choose a noise floor of $N = -85$ dBm in an office scenario with a link distance of $d = 18$ m in order to explore numerically the whole range of SNR while using reasonable values of TXP. The ITU model for indoor attenuation [5] gives a path loss of $L \approx 85$ dBm. Then, we can use (8) to obtain the expected

Table 2: Linear Regressions

Device	ρ_{tx} model estimates (α_i)				ρ_{rx} model estimates (β_i)		
	(Intercept) [W]	MCS [Mbps]	TXP [mW]	adj. r^2	(Intercept) [W]	MCS [Mbps]	adj. r^2
HTC Legend	0.354(14)	0.0052(3)	0.021(3)	0.97	0.013(3)	0.00643(11)	>0.99
Linksys WRT54G	0.540(12)	0.0028(2)	0.075(3)	0.98	0.14(2)	0.0130(7)	0.96
Raspberry Pi	0.478(19)	0.0008(4)	0.044(5)	0.88	-0.0062(14)	0.00146(5)	0.98
Galaxy Note 10.1	0.572(4)	0.0017(1)	0.0105(9)	0.98	0.0409(10)	0.00173(4)	0.99
Soekris net4826-48	0.17(3)	0.0170(6)	0.101(7)	0.99	0.010(8)	0.0237(3)	>0.99

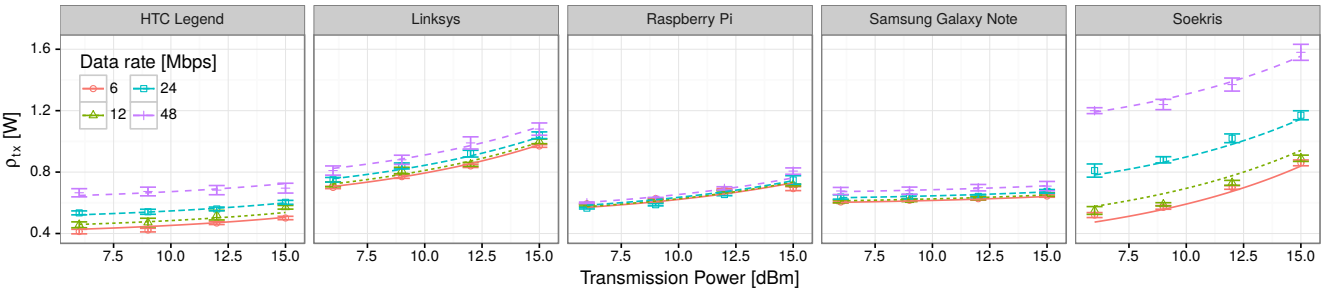

(a) ρ_{tx} fit as a function of MCS and TXP.

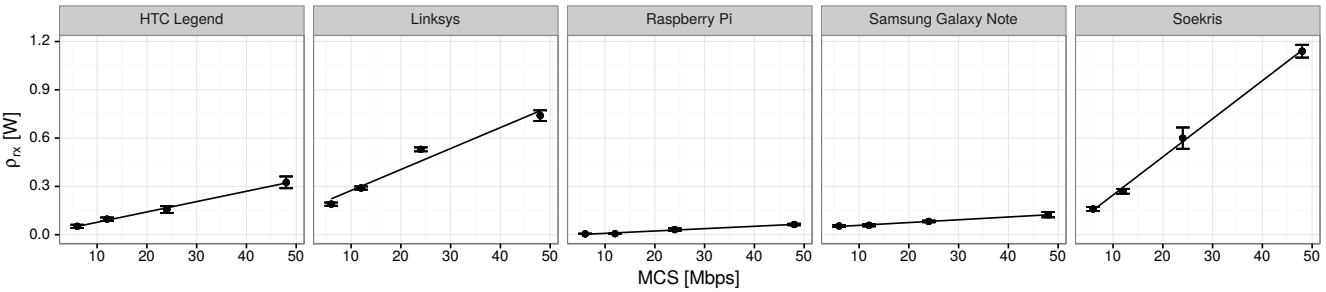

(b) ρ_{rx} fit as a function of MCS.

Figure 2: Linear regressions.

energy consumed per frame and MCS mode, with TXP being directly related to the SNR level.

The results are reported in Fig. 3. As the figure illustrates, consumption first falls abruptly as the TXP increases for all modes, which is caused when the SNR reaches a sharp threshold level such that the number of retransmissions changes from 6 to 0 (i.e., no frame is discarded). From this threshold on, the consumption increases with TXP because, although the number of retransmissions is 0, the wireless interface consumes more power. We note that the actual value of the TXP when the consumption drops depends on the specifics of the scenario considered, but the qualitative conclusions hold for a variety of scenarios.

3.4 Energy Efficiency vs. Optimal Goodput

We can finally merge previous numerical analyses and confront energy efficiency, given by (9), and optimal goodput, given by (10), for all devices and under the aforementioned assumptions. To this aim, we plot in the same figure the energy efficiency for the configuration that maximises goodput given an SNR value vs. the obtained goodput, with the results being depicted in Fig. 4. We next discuss the main findings from the figure.

First of all, we can see that the energy efficiency grows sub-linearly with the optimal goodput (the optimal goodput for each SNR value) in all cases. We may distinguish three different cases in terms of energy efficiency: high (Samsung Galaxy Note and HTC Legend), medium (Raspberry Pi) and low energy efficiency (Linksys and Soekris). Furthermore, for the case of the Soekris, we note that the "central modes" (namely, 4 and 5) are more efficient in their optimal region than the subsequent ones.

Another finding (more relevant perhaps) is that it becomes evident that increasing the goodput does not always improve the energy efficiency: there are more or less drastic leaps, depending on the device, between mode transitions. From the transmitter point of view, in the described scenario, this can be read as follows: we may increase the TXP to increase the SNR, but if the optimal goodput entails a mode transition, the energy efficiency may be affected.

As a conclusion, we have demonstrated that optimal goodput and energy efficiency do not go hand in hand, even in a single spatial stream, in 802.11. There is a trade-off in some circumstances that current rate adaptation algorithms cannot take into account, as they are oblivious to the energy consumption characteristic of the device.

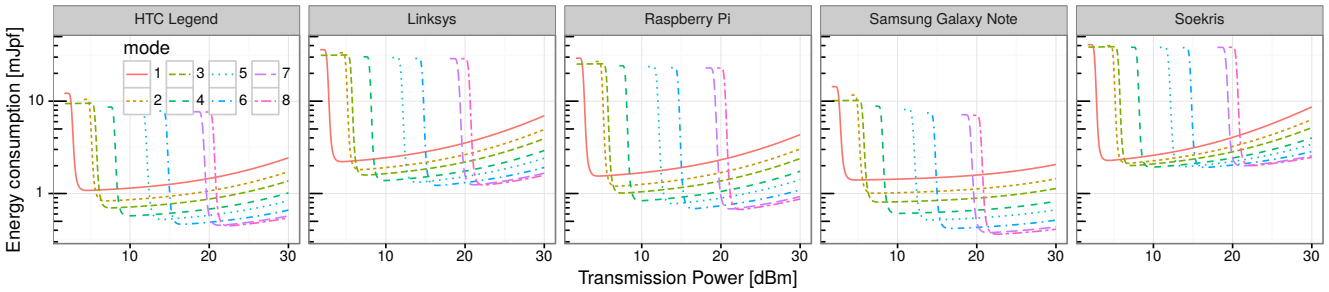

Figure 3: Expected energy consumption per frame in *millijoules per frame* (mJpf) under fixed channel conditions.

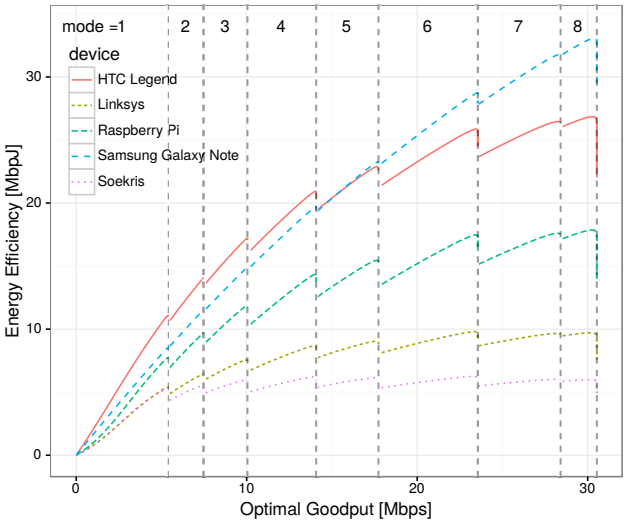

Figure 4: Energy Efficiency vs. Optimal Goodput under fixed channel conditions.

3.5 Sensitivity to Energy Parameter Scaling

We next explore how the different energy parameters affect the energy efficiency vs. optimal goodput relationship. For this purpose, we selected the Raspberry Pi curve from Fig. 4 (results are analogous with the other devices) and we scale up and down, and one at a time, the four energy parameters ρ_{id}, ρ_{tx}, ρ_{rx}, and γ_{xg}. The scaling up and down is done by multiplying and dividing by 3, respectively, the numerical value of the considered parameter. One of the first results is that the impact of ρ_{rx} is negligible, a result somehow expected as the cost of receiving the ACK is practically zero. From this point on, we do not consider further this parameter.

We show in Fig. 5a the overall effect of this parameter scaling. The solid line represents the base case with no scaling (same curve as in Fig. 4), and in dashed and dotted lines the corresponding parameter was multiplied or divided by a factor of 3, respectively. As expected, larger parameters contribute to lower the overall energy efficiency. However, the impact on the energy efficiency drops between mode transitions is far from being obvious, as in some cases transitions are more subtle while in others they become more abrupt.

To delve into these transitions, we illustrate in Fig. 5b the "drop" in energy efficiency when changing between modes.

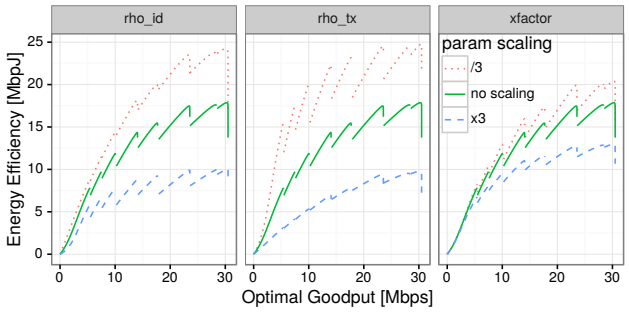

(a) Overall effect.

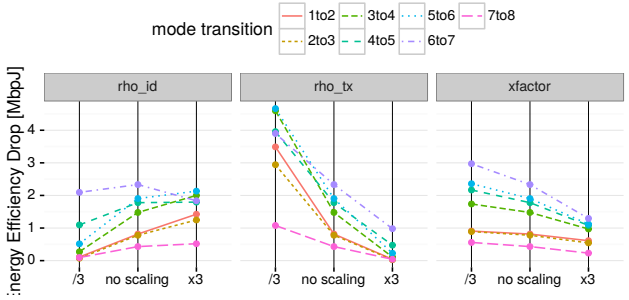

(b) Impact on mode transitions.

Figure 5: Impact of energy parameter scaling on the energy efficiency.

As it can be seen, the cross-factor is the less sensitive parameter of the three, because its overall effect is limited and, more importantly, it scales all the leaps between mode transitions homogeneously. This means that a higher or lower cross-factor, which resides almost entirely in the device and not in the wireless card, does not alter the energy efficiency vs. optimal goodput relationship (note that this parameter results in a constant term in (8)). Thus, the cross-factor is not relevant from the RA point of view, and energy-aware RA algorithms can be implemented by leveraging energy parameters local to the wireless card.

On the other hand, ρ_{id} and ρ_{tx} have a larger overall effect, plus an inhomogeneous and, in general, opposite impact on mode transitions. While a larger ρ_{id} contributes to larger leaps, for the case of ρ_{tx}, the larger energy efficiency drops occur with smaller values of that parameter. Still, the reason

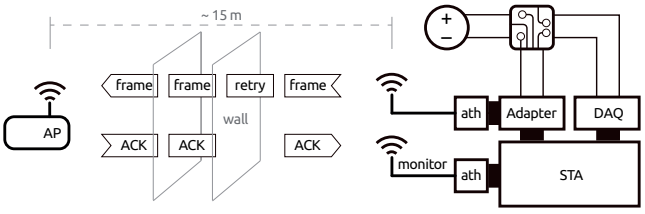

Figure 6: Experimental setup.

behind this behavior is the same for both cases: the wireless card spends more time in idle (and less time transmitting) when a transition to the next mode occurs, which has a higher data rate.

This effect is also evident if we compare the Samsung Galaxy Note and the HTC Legend curves in Fig. 4. Both devices have ρ_{id} and ρ_{tx} in the same order of magnitude, but the HTC Legend has a larger ρ_{id} and a smaller ρ_{tx}. The combined outcome is a more dramatic sub-linear behaviour and an increased energy efficiency drop between mode transitions.

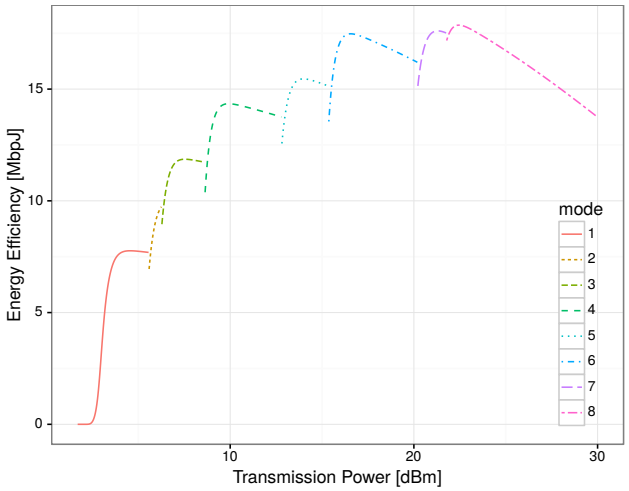

Figure 7: Energy Efficiency vs. Transmission Power under fixed channel conditions for the Raspberry Pi case.

4. EXPERIMENTAL VALIDATION

This section is devoted to experimentally validate the results from the numerical analysis and, therefore, the resulting conclusions. To this aim, we describe our experimental setup and the validation procedure, first specifying the methodology and then the results achieved.

4.1 Experimental Setup

We deployed the testbed illustrated in Fig. 6, which consists of a station (STA) transmitting evenly-spaced maximum-sized UDP packets to an access point (AP). The AP is an x86-based Alix6f2 board with a Mini PCI Qualcomm Atheros AR9220 wireless network adapter, running Voyage Linux with kernel version 3.16.7 and the `ath9k` driver. The STA is a desktop PC with a Mini PCI Express Qualcomm Atheros QCA9880 wireless network adapter, running Fedora Linux 23 with kernel version 4.2.5 and the `ath10k` driver. We also installed at the STA a Mini PCI Qualcomm Atheros AR9220 wireless network adapter to monitor the wireless channel.

The QCA9880 card is connected to the PC through a *x1 PCI Express to Mini PCI Express* adapter from Amfeltec. This adapter connects the PCI bus' data channels to the host and provides an ATX port so that the wireless card can be supplied by an external power source. The power supply is a Keithley 2304A DC Power Supply, and it powers the wireless card through an *ad-hoc* measurement circuit that extracts the voltage and converts the current with a high-precision sensing resistor and amplifier. These signals are measured using a National Instruments PCI-6289 multifunction data acquisition (DAQ) device, which is also connected to the STA. Thanks to this configuration, the STA can simultaneously measure the instant power consumed by the QCA9880 card[1] and the goodput achieved.

As the figure illustrates, the STA is located in an office space and the AP is placed in a laboratory 15 m away, and transmitted frames have to transverse two thin brick walls. The wireless card uses only one antenna and a practically-empty channel in the 5-GHz band. Throughout the experiments, the STA is constantly backlogged with data to send to the AP, and measures the throughput obtained by counting the number of received acknowledgements (ACKs).

4.2 Methodology and Results

In order to validate our results, our aim is to replicate the qualitative behaviour of Fig. 4, in which there are energy efficiency "drops" as the optimal goodput increases. However, in our experimental setting, channel conditions are far from steady, which introduces a notable variability in the results as it affects both the x-axis (goodput) and the y-axis (energy efficiency). To reduce the impact of this variability, we decided to change the variable in the x-axis from the optimal goodput to the transmission power –a variable that is directly configured in the wireless card. In this way, the qualitative behaviour to replicate is the one illustrated in Fig. 7, where we can still identify the performance "drops" causing the loss in energy efficiency.

Building on this figure, we perform a sweep through all available combinations of MCS (see Table 1) and TXP[2]. Furthermore, we also tested two different configurations of the AP's TXP at different times of the day, to confirm that this qualitative behaviour is still present under different channel conditions. For each configuration, we performed 2-second experiments in which we measure the total bytes successfully sent and the energy consumed by the QCA9880 card with sub-microsecond precision. From this data, the energy efficiency is computed, with the results depicted in Fig. 8 (each figure corresponds to a different TXP value configured at the AP).

In the figure, each line type represents the STA's mode that achieved the highest goodput for each TXP interval, therefore in some cases low modes do not appear because

[1]Following the discussion on Section 3.5, the device's cross-factor is not involved in the trade-off, thus we will expect to reproduce it by measuring the wireless interface alone.

[2]The model explores a range between 0 and 30 dBm to get the big picture, but this particular wireless card only allows us to sweep from 0 to 20 dBm.

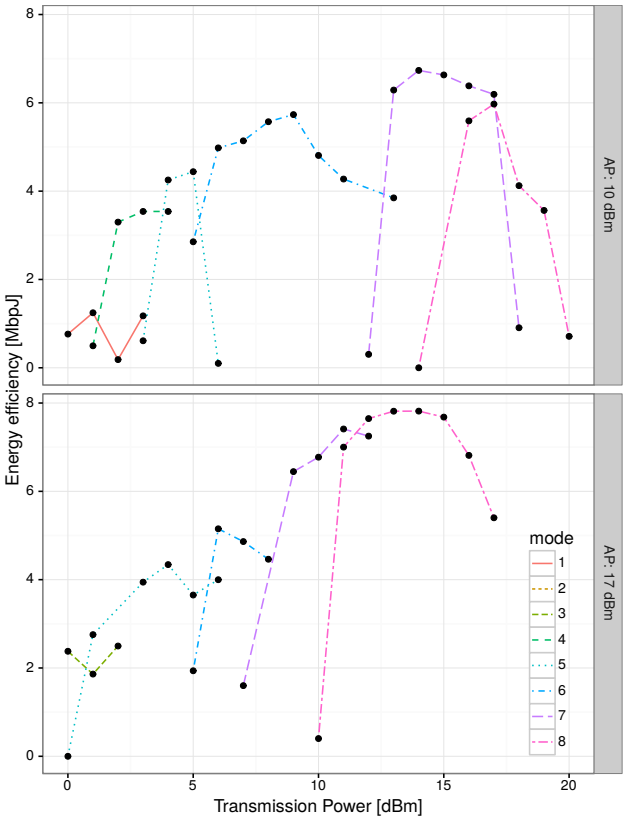

Figure 8: Experimental study of Fig. 7 for two AP configurations.

a higher mode achieved a higher goodput. Despite the inherent experimental difficulties, namely, the low granularity of 1-dBm steps and the random variability of the channel, the experimental results validate the analytical ones, as the qualitative behaviour of both figures follows the one illustrated in Fig. 7. In particular, the performance "drops" of each dominant mode can be clearly observed (especially for the 36, 48 and 54 Mbps MCSs) despite the variability in the results.

5. SUMMARY AND FUTURE WORK

In this paper, we have revisited 802.11 rate adaptation by taking energy consumption into account. While some previous studies pointed out that MIMO rate adaptation is not energy efficient, we have demonstrated through numerical analysis that, even for single spatial streams without interfering traffic, energy consumption and throughput performance are different optimisation objectives. Furthermore, we have validated our results via experimentation.

Our findings show that this trade-off emerges at certain "mode transitions" when maximising the goodput, suggesting that small goodput degradations may lead to energy efficiency gains. Moreover, our analyses have showed that these trade-offs arise as a consequence of the power consumption behaviour of wireless cards and does not depend on the energy consumed in the rest of the device. In this way, energy-aware rate adaptation may be achieved building on information local to the wireless interface. Still, to develop energy-aware rate adaptation algorithms, further research is needed to understand how the findings of this work can be leveraged in suboptimal conditions, and how other effects, such as collisions and MIMO, affect the established trade-off.

6. ACKNOWLEDGEMENTS

This work has been performed in the framework of the H2020-ICT-2014-2 projects 5GNORMA (grant agreement no. 671584) and Flex5Gware (grant agreement no. 671563). The authors would like to acknowledge the contributions of their colleagues. This information reflects the consortium's view, but the consortium is not liable for any use that may be made of any of the information contained therein.

7. REFERENCES

[1] Cisco. Cisco visual networking index: forecast and methodology, 2014-2019. Technical report, Cisco, 2015.

[2] S. Eryigit, G. Gur, S. Bayhan, and T. Tugcu. Energy efficiency is a subtle concept: fundamental trade-offs for cognitive radio networks. *Communications Magazine, IEEE*, 52(7):30–36, July 2014.

[3] A. Garcia-Saavedra, P. Serrano, A. Banchs, and M. Hollick. Balancing energy efficiency and throughput fairness in IEEE 802.11 WLANs. *Pervasive and Mobile Computing*, 8(5):631 – 645, 2012.

[4] K. Huang, K. Duffy, and D. Malone. H-RCA: 802.11 Collision-Aware Rate Control. *Networking, IEEE/ACM Transactions on*, 21(4):1021–1034, Aug 2013.

[5] ITU-R. Recommendation P.1238, International Telecommunication Union, July, 2015.

[6] M. O. Khan, V. Dave, Y.-C. Chen, O. Jensen, L. Qiu, A. Bhartia, and S. Rallapalli. Model-driven energy-aware rate adaptation. In *Proceedings of the Fourteenth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, MobiHoc '13, pages 217–226, New York, NY, USA, 2013. ACM.

[7] C.-Y. Li, C. Peng, S. Lu, and X. Wang. Energy-based Rate Adaptation for 802.11n. In *Proceedings of the 18th Annual International Conference on Mobile Computing and Networking*, Mobicom '12, pages 341–352, New York, NY, USA, 2012. ACM.

[8] D. Qiao, S. Choi, and K. Shin. Goodput analysis and link adaptation for IEEE 802.11a wireless LANs. *Mobile Computing, IEEE Transactions on*, 1(4):278–292, Oct 2002.

[9] P. Serrano, A. de la Oliva, P. Patras, V. Mancuso, and A. Banchs. Greening wireless communications: Status and future directions. *Computer Communications*, 35(14):1651 – 1661, 2012. Special issue: Wireless Green Communications and Networking.

[10] P. Serrano, A. Garcia-Saavedra, G. Bianchi, A. Banchs, and A. Azcorra. Per-Frame Energy Consumption in 802.11 Devices and Its Implication on Modeling and Design. *IEEE/ACM Transactions on Networking*, PP(99):1–1, 2014.

[11] I. Ucar and A. Azcorra. Deseeding Energy Consumption of Network Stacks. In *Research and Technologies for Society and Industry Leveraging a better tomorrow (RTSI), 2015 IEEE 1st International Forum on*, pages 7–16, Sept 2015.